Flattening generalization gradients, context, and perceptual learning

JAMES BYRON NELSON AND MARIA DEL CARMEN SANJUAN
Universidad del País Vasco, San Sebastián, Gipuzkoa, Spain

The present research investigated the effects of physical context change and perceptual learning on generalization. In a video game, participants learned to suppress their mouse-clicking behavior in the presence of one stimulus (AX). Generalization was observed between the AX stimulus and another stimulus (BX) that was designed to be similar. When testing was conducted in a context different from that in which AX was used in training, responding to AX was attenuated, and responding to BX was enhanced. That is, the generalization gradient flattened. The latter effect was only evident in groups for which generalization had been reduced through a preexposure manipulation believed to produce perceptual learning. Experiment 2 demonstrated that the increase in generalization observed in the first experiment was due to the context change between the preexposure and test rather than to a change between the conditioning and test contexts. Implications for flattening generalization gradients and mechanisms of perceptual learning are discussed.

The present work was undertaken to assess whether a change in a physical background context could affect generalization between similar stimuli. Generalization is a well-known construct evoked by theorists when similar responding is observed to a stimulus involved in some type of training and to a different test stimulus. The interesting observation about generalization that concerned the present work is that it is not static. Generalization can sometimes appear to increase over time.

The increase in generalization that is sometimes observed with the passage of time is referred to as a "flattening" of the generalization gradient (see, e.g., Riccio, Richardson, & Ebner, 1984). The responding to the trained target stimulus and test stimuli becomes more similar. Bouton, Nelson, and Rosas (1999) noted that such flattening can be observed in two ways. First, there could be a loss of responding to the target stimulus—perhaps as the result of its being perceived as more similar to the test stimulus (which evokes little responding on its own). Second, there could be an increase in responding to the test stimuli—perhaps as the result of their being perceived as more similar to the target stimulus. Flattening of the generalization gradient often occurs by combinations of both types of change in responding (see Bouton et al., 1999, for a review).

To explain flattening generalization gradients, Bouton et al. (1999) suggested that such effects could be caused by contextual change leading to forgetting of specific stimulus attributes. Assuming that time provides a type of context for events embedded within it, the passage of time should cause some forgetting of the specifics of those events, leading to an increase in generalization. That explanation leads to the question as to whether a change

in context provided by physical background cues would produce such flattening. Assessing that question was the primary goal of the present research.

To address the question, we used human participants in a video-game preparation (see Figure 1) that Nelson and Sanjuan (2006) used to demonstrate latent inhibition. A baseline of computer-mouse clicking was established in which participants clicked a computer mouse that launched torpedoes at an on-screen spacecraft. Stimuli were presented in the form of colored oval "sensors" near the bottom of the screen, which was constructed so that the majority of the sensor appeared as green, with either a small blue or yellow portion to its left. These stimuli were referred to as *AX* and *BX*, with X indicating the green portion of the stimuli that was common, and A and B referencing the unique blue and yellow components. Different contexts were provided by different "space" backgrounds that were shown through the viewport and upon which the action in the game took place.

To assess generalization, the AX stimulus was paired with an attack unconditioned stimulus (US) from the on-screen spacecraft. Participants were instructed that they should suppress their own rates of torpedo launching to conserve energy when that they believed they were about to be attacked. Thus, as participants learned to expect an unavoidable attack following the AX stimulus, they suppressed their own mouse clicking. Following conditioning of AX, generalization was assessed by presenting the BX stimulus four times in extinction. Suppression to BX should vary as a function of its perceived similarity to AX.

Bouton et al. (1999) pointed out that changes in the generalization gradient over time are not necessarily very

J. B. Nelson, drjbn@hotmail.com

Figure 1. Video game screen used in Experiments 1–2. Stimuli were presented in the middle oval at the bottom. The white portion was colored green and the small black portion left of the white was colored yellow or blue. The background upon which the spaceship appeared served as the context. The oval appeared black when no stimuli were present.

large effects, and most demonstrations tend to use stimuli that are likely to be perceived as very similar in the absence of flattening. Butler (1964; discussed in Desiderato, Butler, & Meyer, 1966) suggested that the change between the training and test stimuli must be moderate to observe changes in generalization over time. Later work has been largely consistent with that idea. For example, Thomas et al. (1985) trained two groups of pigeons in a line discrimination task in which a line oriented at 0º served as S+. In one group, the S− was a line oriented at 90º, and in the other, the S− was a line oriented at 30º. When tested immediately, the groups did not differ, and approximately 94% of their responses were directed toward the S+. When tested after 1 week, the groups did differ. Those trained with a 90º S− directed about 88% of their responses to the S+, and those trained with the 30º S− directed only about 72% of their responses to the S+. With the passage of time, more responses came to be directed at the similar S− than at the dissimilar S−. The perceptual difference between the S+ and S− stimuli was arguably smaller in the latter group, and the change in their discriminability over time was greater (see Kraemer, 1984; Zhou & Riccio, 1996, for similar results)

The observation that the flattening of a generalization gradient is most reliably observed with physically similar stimuli works against the ability to detect a change in that gradient. Optimal conditions under which to detect such an effect would be where the generalization gradient was sharp enough to provide room to observe a change. Thus, we included one group that underwent a perceptual-learning manipulation to provide potentially optimum conditions to observe an effect of context change.

Perceptual learning is said to narrow the gradient of generalization. That is, animals and people respond as if they perceive the training and test stimuli more differently as the result of manipulations thought to produce percep-

tual learning (see Goldstone, 1998; Hall, 1991, 2003; McLaren & Mackintosh, 2000, for reviews). Perhaps the most potent manipulation to produce perceptual learning involves exposure to the relevant stimuli in alternation (see, e.g., Lavis & Mitchell, 2006; Symonds & Hall, 1995). Following that precedent, one group in the present experiments received alternating preexposures (AX/BX/AX/BX . . .) to the stimuli prior to conditioning and testing. This manipulation was designed to reduce the generalization between the stimuli without modifying their physical properties, perhaps allowing for greater potential to observe an effect of the context change.

EXPERIMENTS 1A AND 1B

The first set of studies was undertaken to determine whether a change in context between training and testing would lead to a pattern of responding indicative of a flattening of the generalization gradient. In Experiment 1A, we assessed generalization to BX in either the same context as that in which AX was trained, or in a different context. In Experiment 1B, we assessed the impact of a change of context on responding to AX. When testing takes place in a different context from that in which training took place, an apparent flattening of the generalization gradient could be observed as an increase in suppression to BX, suggesting an increase in generalization to the test stimulus. Flattening could also be observed as a decrease in suppression to the training stimulus, suggesting an increase in generalization from the training to the test stimuli.

The design of the present experiment is shown in the top portion of Table 1. Participants in both Experiments 1A and 1B received simple conditioning with AX either without any pretraining (Group None), or after a perceptual learning manipulation prior to conditioning (Group Intermixed). In the latter condition, participants received four alternating preexposures to the AX and BX stimuli. Following conditioning, participants from each condition received a test either in the context in which preexposure and conditioning took place, or in a different context. Participants in Experiment 1A received testing with the BX stimulus, and participants in Experiment 1B received testing with the AX stimulus used in conditioning. Otherwise, the treatments in the two studies were identical. The physical difference between AX and BX was small. Thus, there should have been substantial generalization between them, and there may have been little room to observe any change in that generalization. In Group Intermixed, the preexposure to the stimuli prior to conditioning should decrease generalization, allowing for more room to observe an effect of the test context.

Method

Participants
Forty-three volunteers participated in Experiment 1A, and 41 participated in Experiment 1B.

Apparatus
Since we used the exact same apparatus as that in Nelson and Sanjuan (2006), we will provide a summary description here. Participants

Table 1
Design of Experiments

Group	Preexposure	Conditioning	Test 1A	Test 1B
Experiments 1A and 1B				
Intermixed	A:AX/BX		A:BX−	B:AX−
		A:AX+	or	or
None	A:_____		B:BX−	B:AX−
Experiment 2				
Intermixed	A:AX/BX	A:AX+ or B:AX+	A:BX− or B:BX−	
None	A:_____			

Note—Preexposure indicates the stimuli to which the participants were preexposed prior to conditioning. A: and B: are different contexts, whereas AX and BX are different compound stimuli composed of yellow or blue (A or B, counterbalanced) and green (X). A "+" indicates a pairing with an attack US, and the absence of the "+" or a "−" indicates no such pairing. In Experiment 2, the lines connecting across phases indicate how each group was divided across phases. To illustrate, after preexposure, Groups Intermixed and None were each divided, with half of the participants from each condition in context A: or B: (see text for details).

received written instructions that they were playing a game in which they were to earn points by shooting torpedos at an on-screen spaceship by clicking the mouse. They were further instructed that sometimes they would be attacked and that the attack would damage their spaceships by draining their power, leaving the participants unable to continue the game until power was recharged. They were instructed that they could not avoid the attack, but that they could prepare for it by conserving their power (suppressing their own rate of torpedo firing) when they believed they were about to be attacked. They were told that sensors would appear that might help them in the game, and they were not told what the sensors would indicate. The instructions were the same as those reported in detail in Nelson and Sanjuan.

The video game (see Figure 1) was viewed on a standard 15-in. (38.1-cm) computer monitor, with a resolution of 800 × 600 pixels and 24-bit color depth. On the monitor, an image was presented so that it was as if the participant were sitting inside of a spaceship looking out of a viewscreen. The viewscreen was a rectangular window that was 618 pixels wide and 368 pixels tall (23.5 × 14 cm) and that was centered from left to right, 3.5 cm below the top of the screen on a grey metallic background. A box appeared 1.3 cm below the top of the screen in which the word "Points" appeared in yellow. At the bottom of the screen, five black ovals appeared that were each 3.28 cm in diameter. The third was centered from left to right and was located .8 cm above the bottom of the screen. The other four ovals were spaced at intervals of approximately 2 cm to the left and right of the center oval. One of two colored backgrounds (Hubble Space Telescope photos of the Eagle 1 or Crab Nebulae) could be seen through the viewscreen on which a three-dimensional representation of a spaceship was flying in a randomly determined path. These colored backgrounds provided contexts and were always counterbalanced.

Stimuli were presented as the 5-sec illumination of the middle oval. The middle oval had a diameter of 85 pixels (approx. 3.28 cm on the screen), and a total area of 5,310 pixels. Stimuli were composed of the elements X, A, and B. The X portion was created by illuminating the oval, beginning at the 12th pixel along the diameter from the left, with the color green filling 4,905 of the 5,310 pixels. Elements A and B were created by illuminating the oval along the first 11 pixels along the diameter (405 pixels), with either blue or yellow (A or B, counterbalanced). Stimuli AX and BX were presented by illuminating the portions of the oval corresponding to A or B, and X.

The hypothetical "unconditioned stimulus" was presented in the form of an inescapable attack from the enemy spaceship. Immediately upon the offset of a sensor stimulus, a round green torpedo

emerged from the on-screen spacecraft and exploded in the center of the viewscreen. The message "Power at ___ percent. Controls frozen for ____ seconds." appeared in the center of the viewscreen and remained until "Power" incremented to 100 and "Controls frozen for ____" decremented to 0 (changes occurring roughly every second). During this time, the computer mouse was inoperable, and actions of the participant were not reflected on the screen. The numbers in the blanks were determined by a suppression ratio in which the number of clicks during the 5 sec prior to the attack (i.e., during the presentation of the sensor) was divided by that number plus the average rate of mouse clicks in the 5 sec before stimulus presentations across the game. The resulting ratio was then multiplied by 120. For example, if a participant clicked, on average, 10 times prior to the sensor conditioned stimulus (CS) and did not suppress his or her rate of clicking (clicking 10 times during the sensor CS), then the ratio would calculate to .5, and their controls would be frozen for 60 sec by an attack.

Procedure

The availability of the experiment was announced to potential participants, and they had approximately 1 month to volunteer in each study. Our experience with this method had indicated that a 1-month volunteer period in an academic semester led to substantial variability between the participants who volunteer early in that time period and those who volunteer near the deadline. Nevertheless, such a long period was typically necessary to allow for adequate sample sizes. Thus, we made no attempt to equalize group sizes; doing so had the potential to nonrandomly force participants who volunteered near the end of the deadline into particular groups. Each participant was randomly assigned to a group.

Up to 14 participants could be run at one time, although they were run in groups of 5 or less. They were distributed as widely as possible across space in the room. They were seated at their individual computers, read and signed a consent form, and were given the instructions (exact instructions are available in Nelson & Sanjuan, 2006). They read the instructions to themselves and then had the instructions read to them by the attending researcher. They were instructed to place their left hand on the "s" key and their right hand on the mouse. The lights in the room were turned off, and they were instructed to press the "s" key, starting the game.

Preexposure and baseline. Following the procedure of Nelson and Sanjuan (2006), all participants began by playing the game for 60 sec in the different context (B) to ensure familiarity with that context. During this time—and for the remainder of the game—they clicked the mouse on a variable-ratio 3 schedule in which a random one in three clicks led to the launch of a torpedo. A random half of those torpedoes "locked onto" the enemy spacecraft and followed it until it exploded on the spacecraft, which added a point to the point counter at the top of the screen. The spacecraft was never destroyed; this maintained continuity across the game.

After those 60 sec, a context switch was initiated using the procedure of Nelson and Sanjuan (2006). A standard Microsoft Windows message box appeared in the middle of the screen. The text "Please attend to this important message" was displayed on the title bar of the message box, and the text in the message box read, "You, your sensors, and the enemy are being transported to another galaxy for further testing. Press 'OK' to proceed." When the participant pressed the "OK" button, another message box was displayed. The text "Press 'OK' now for immediate transport" was displayed in the title bar, and the text in the message box read, "Remember, you, your sensors, and the enemy are being transported to another galaxy." When the participant pressed the "OK" button, the screen flickered, and the current background (e.g., Crab Nebula) was replaced with the alternative (e.g., Eagle 1 Nebula).

At that point, preexposure began for Group Intermixed, whereas Group None simply played the game with no events for the same amount of time. Group Intermixed received four nonreinforced 5-sec exposures to the AX and BX stimuli in an alternating fashion, with the order of exposure counterbalanced between subjects. The

intertrial interval (ITI) on these trials averaged 11.2 sec. Treatment of these groups was identical in both Experiments 1A and 1B.

Conditioning of AX. Conditioning consisted of 10 presentations of the AX stimulus, with each presentation of the AX stimulus paired with an attack US. The ITI (US offset to CS onset) averaged 11.2 sec. The first conditioning trial began 11 sec after the last pre-exposure trial.

Testing. Testing was conducted in either the context in which AX was conditioned, or in the different context. For participants being tested in a different context, a context switch was initiated, as was described earlier. Participants in Experiment 1A received four presentations of the BX stimulus in extinction. Participants in Experiment 1B were tested with the AX stimulus. The ITI on these trials averaged 11 sec.

Data Analysis

Data. The number of times the participant clicked the mouse during the 5-sec CS and during the 5 sec preceding the CS was recorded on each trial. Standard suppression ratios [CS / (CS + pre-CS)] were calculated.

Exclusion criteria. Participants sometimes had zero responses in the pre-CS, making an interpretation of the resulting ratio difficult. We adopted the procedure used by Nelson and Sanjuan (2006) of excluding participants who had an average response rate less than 1 click/sec. Rates below 1 click/sec create a very coarse measure of suppression with the 5-sec CS, and this criterion excluded all participants who had zero pre-CS rates on some trials. Because detection of generalization depends on good conditioning, we also excluded participants for whom suppression ratios on two of the last three conditioning trials were .4 or higher. χ^2 tests of independence were used to assess whether the selection criteria were independent of the grouping variable. To anticipate, there was no relationship between exclusion and group membership.

Hypothesis tests. The random assignment procedure left the design slightly unbalanced; there were unequal sample sizes between groups and counterbalance conditions. The selection criteria also contributed to the inequality. Since those criteria were independent of group, the inequality of the sample sizes was random.

Suppression-ratio and pre-CS data were analyzed with a mixed-factorial ANOVA using the typical Type III (unique) sums of squares. This unweighted means approach—combined with the fact that the varying sample size was a random factor—prevents interpretation from being complicated by the slight unbalance.

Simple-effect tests were conducted with an ANOVA using error terms derived from the overall analysis. Degrees of freedom were reduced using the Welch (1938)–Satterthwaite (1946) procedure to compensate for the pooling of potentially heterogeneous variances. Throughout, a rejection criterion of $p < .05$ was adopted, although exact probabilities are reported for the reader.

Results

Exclusion Criteria

Experiment 1A. Random assignment placed 23 participants in Group Intermixed, with 11 tested in context same, and 12 tested in context different. There were 20 participants in Group None, with 10 tested in each context. Application of the exclusion criteria removed 2 participants from Group Intermixed—1 from each test condition. Three participants were removed from Group None, 2 coming from those tested in context same, and 1 from those tested in context different. Exclusion was independent of membership in the four conditions ($\chi^2 < 1$).

Experiment 1B. Random assignment placed 24 participants in Group Intermixed, with 12 tested in each context. There were 17 participants in Group None, with 9 tested in context same and 8 tested in context different.

The application of the exclusion criteria removed 4 participants from Group Intermixed, with 3 coming from those tested in context same and 1 from those tested in context different. Four participants were removed from Group None; 1 came from those tested in context same, and 3 from those tested in context different. Exclusion was independent of membership in the four conditions ($\chi^2 = 3.23$, $p = .36$).

Conditioning

Suppression ratios. Training was identical for Experiments 1A and 1B, and the data were combined for the analysis with an exposure (intermixed or none) × context (to be tested in context same or context different) × test type (to be tested with BX or AX) × trials ANOVA. There was an effect of trials [$F(9,567) = 46.49$, $p < .001$], a main effect of exposure [$F(1,63) = 5.17$, $p = .03$], and an exposure × trials interaction [$F(9,567) = 4.18$, $p < .001$]. There were no other effects in the overall analysis ($ps \geq .29$). The data are shown at left in Figure 2, which collapses across the test type variable and the context variables.

Simple effect tests showed that the two exposure conditions differed on Trial 2 [$F(1,348) = 39.03$, $p < .001$], with less suppression in Group Intermixed than in Group None, but not on any other trial ($Fs \leq 3.47$, $ps \geq .06$). Of importance for the discussion to follow, the groups did not differ on Trial 3 or 10 ($F < 1$), yet suppression in both groups was less on Trial 10 than on Trial 3 ($Fs \geq 42.35$, $ps \leq .001$). Thus, there was room in the response scale to see differences after Trial 2, if such differences existed.

Pre-CS. The same analysis applied to the pre-CS response counts showed an effect of trials [$F(9,567) = 2.06$, $p = .03$]. Pre-CS responding increased from an average of 16.21 ($SD = 6.35$) on Trial 1 to 17.08 ($SD = 7.86$) on Trial 10. There were no other effects in the analysis of the pre-CS responding ($ps \geq .08$).

Experiment 1A Test

Suppression ratios. An exposure × context × trials analysis of all four test trials showed a reliable effect of trials [$F(3,102) = 26.15$, $p < .001$], an effect of exposure [$F(1,34) = 6.55$, $p = .01$], and an overall exposure × context interaction [$F(1,34) = 7.32$, $p = .01$]. No other effects were reliable, although the trials × context interaction approached reliability [$F(3,34) = 2.31$, $p = .08$].

Simple-effect tests of the average suppression to BX investigating the exposure × context interaction showed an effect of context in Group Intermixed where there was less suppression in Group Intermixed Same than in the other three groups [$Fs(1,34) \geq 5.34$, $ps \leq .03$]. Group Intermixed Different did not differ from Groups Novel Same or Novel Different [$Fs(1,34) \leq 1.72$, $ps \geq .19$]. The difference between Groups Novel Same and Novel Different was not reliable [$F(1,34) = 1.83$, $p = .18$]. Changing the context led to a substantial increase in suppression in the group receiving intermixed exposure.

Additional comparisons were made with Group None. First, a trials × context ANOVA was conducted, comparing the last trial of conditioning with the first trial of testing. That analysis showed an effect of trials [$F(1,15) = 6.23$,
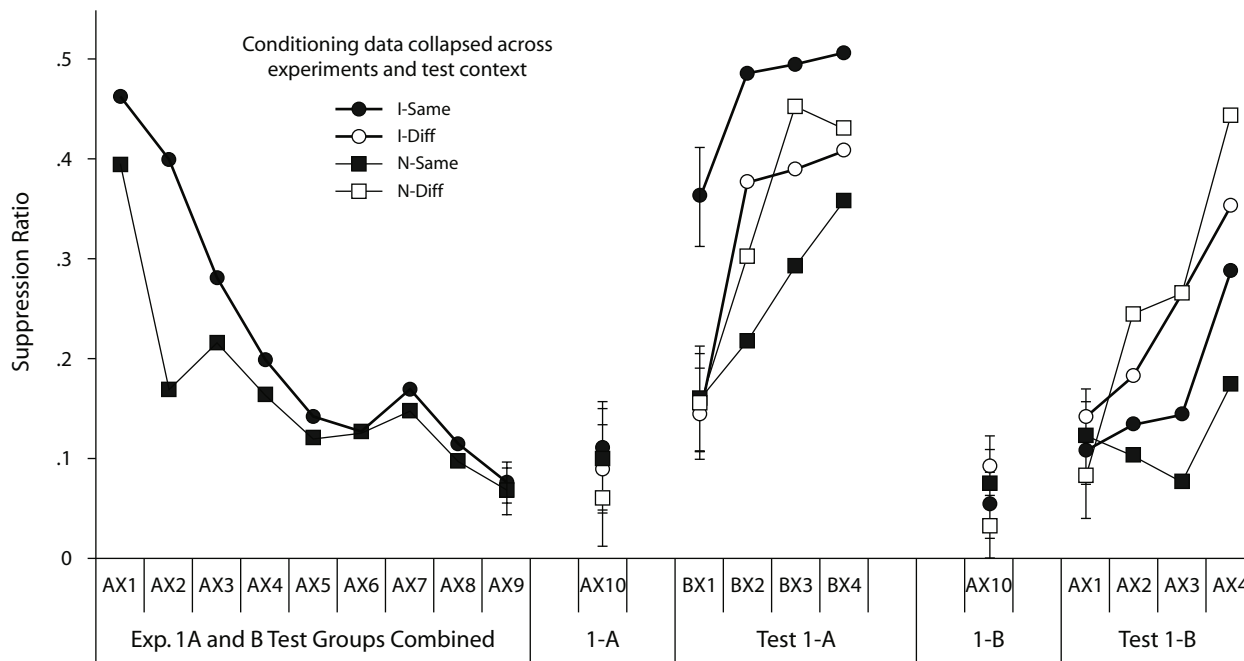
Figure 2. Data are shown from Experiments 1A and 1B. The acquisition of suppression to AX from both experiments following intermixed (I) preexposure to AX and BX, or none (N), is shown in the first 9 trials at left, collapsed across group and test context. The middle set of points shows suppression to AX on Trial 10 in Experiment 1A and generalization testing with BX in either the same context as preexposure and conditioning or a different context. The right set of points shows suppression to AX on Trial 10 in Experiment 1B and generalization testing with AX in either the same context as preexposure and conditioning or a different context. Error bars indicate the standard error of the mean.

$p = .02$] and no other effects ($Fs < 1$). Thus, generalization from AX to BX was not complete at the start of the test. The main effects reported above were on the average responding on the test, and some extinction of suppression was reflected in that measure. Comparisons of suppression to AX on the last trial in the group tested in the same context with the average response to BX on the test in the group tested in the different context, or vice versa, were both reliable [$Fs(1,15) \geq 5.95$, $ps < .03$]. Although these analyses obviously confound group with context, they were intended only to demonstrate that with the average level of suppression observed in Group None, between-subjects differences could be detected. It was unlikely that the lack of effect of context in these groups was an artifact of a response floor or complete generalization.

Pre-CS. An exposure × context × trials analysis of pre-CS responding on all four test trials found no effects ($Fs \leq 2.03$, $ps \geq .11$). Pre-CS responding averaged 19.03 ($SD = 5.16$) in this phase.

Experiment 1B Test

Suppression ratios. Data from the last conditioning trial with AX and the test in Experiment 1B are shown at the right in Figure 2. An exposure × context × trials ANOVA of all four test trials showed main effects of trials [$F(3,87) = 9.21$, $p < .001$] and context [$F(1,29) = 6.28$, $p = .018$], where suppression was overall less in context different than in context same. No other effects were reliable ($Fs < 2.12$, $ps > .10$).

Planned simple-effect tests largely corroborated the overall analysis. The groups did not differ within either context ($Fs < 1$). However, the simple effect of context was reliable in Group No Exposure [$F(1,29) = 4.63$, $p = .04$], but not in Group Intermixed [$F(1,29) = 1.74$, $p = .20$].

Pre-CS. An exposure × context × trials analysis of the pre-CS response counts on the test produced no effects ($ps > .20$). Pre-CS responding averaged 16.48 ($SD = 7.23$) in the test phase.

Discussion

Participants received either no preexposure or intermixed preexposure to stimuli AX and BX prior to the conditioning of AX. They then received a test in either the context in which preexposure and conditioning took place, or in a different context. Participants in Experiment 1A were tested with BX, and those in Experiment 1B were tested with AX. A change in context produced a flattening of the generalization gradient so that responding to the training and test stimuli became more similar. Responding to AX was attenuated with a context switch, and responding to BX was enhanced substantially with a context switch, but only when it had been involved in perceptual learning.

In these studies, there was a reliable effect of exposure on conditioning, where conditioning to AX proceeded slightly more slowly following the intermixed AX/BX exposure than it did when the stimulus was novel. This effect is consistent with the idea that latent inhibition accrued to AX (see Nelson & Sanjuan, 2006, for a demonstration of

latent inhibition with this method). Although Group Intermixed showed slightly slower acquisition of suppression than Group None, the groups converged on Trial 3, and suppression continued to increase reliably across a measurable range. Thus, it is unlikely that the groups differed in their terminal levels of conditioning to AX.

There was very little generalization to BX following intermixed preexposure when testing was in the same context as that used for preexposure and conditioning. When testing of BX was in a different context, the reduction in generalization afforded by the preexposure was completely lost. In Group None, there was no detectable effect of changing the context. Despite the trend for less suppression to BX when tested in a different context, generalization to BX was statistically the same for groups that had received no preexposure, whether tested in the same or a different context.

When testing occurred with AX, there was a main effect of context that reflected that suppression elicited by AX was generally less when testing was in a different context. Although the simple-effect tests found an effect in only one group, to conclude that there was no effect in Group Intermixed would likely be an error. That suggestion is strengthened by the lack of a statistically reliable group × context interaction in the overall analysis and the decrease in power that accompanies the reduction in N between the main effect and the simple-effect tests. The most likely conclusion is that simple conditioning is generally context specific with this conditioning method. However, the mechanism behind that context specificity is unclear. It could result from a simple generalization decrement in the perception of the CS or by some type of hierarchical control of the CS–US association by the context.

The effect observed in the groups that received preexposure is consistent with two other published reports that indicate that contextual change can disrupt perceptual learning. Petrov, Dosher, and Lu (2006) have shown that changes in background stimulation can cause significant disruptions in sensitivity (d') to detect subtle changes in stimuli. Participants were required to discriminate between two differently oriented visual intensity maps (produced by a Gabor function) presented on different noisy visual backgrounds. The training was extensive with 300 trials/block and four blocks/day, over 9 days. Training was arranged so that the visual-noise background changed in the middle of a block of trials every eight blocks. From the beginning to the end of training, these changes in background stimulation consistently caused a transient reduction in participants' ability to discriminate the two orientations.

Similar results were obtained by Crist, Kapadia, Westheimer, and Gilbert (1997). In a visual bisection task, participants observed three small lines and were required to determine whether the center line was closer to the left or right line. Accuracy improved substantially over the course of training. Yet, the effects of that training did not generalize when the task was switched to a vernier discrimination, although the location of the lines in the visual field was the same. A simple change in the arrangement of the stimuli led to a loss of discriminative ability.

In sum, these experiments demonstrate that a change in a physical context can produce behavior indicative of some flattening of the generalization gradient. When generalization has been reduced through intermixed preexposure to stimuli, the effect was more pronounced. Simple conditioning was also shown to be somewhat context specific, making for an interesting interaction in the present findings. After intermixed exposure, a change in context tended to increase generalization, whereas the general effect of a context change was to decrease associatively based responding. Thus, generalization that has been reduced by preexposure appeared to increase in the new context to the extent that the initial learning also transferred.

EXPERIMENT 2

Experiments 1A and 1B showed an effect of context change on both perceptual learning and simple conditioning when both the preexposure and conditioning contexts were different from the test context. Perceptual learning is often thought to involve latent inhibition (see Hall, 2003; Lubow, 1989; and McLaren & Mackintosh, 2000, for reviews). When participants are exposed to the AX and BX stimuli, they receive twice as much exposure to the X element; thus, this element should accrue more latent inhibition and be less able to enter into associations than the unique elements. Latent inhibition is also known to be context specific (see, e.g., Nelson & Sanjuan, 2006). When preexposure and conditioning are conducted in different contexts, stimuli enter into associations more normally, as if they had not been preexposed.

Previous studies in animal literature have investigated the effects of contextual change on perceptual learning when that change has been introduced between preexposure and conditioning (Channell & Hall, 1981; Lubow, Rifkin, & Alek, 1976; Trobalon, Chamizo, & Mackintosh, 1992). The findings have been mixed. Those reported by Channell and Hall—as well as those by Lubow et al.—appeared to indicate that perceptual learning was more apparent when the preexposed stimuli were encountered in a new context. The experiments by Trobalon et al., which controlled for the novelty of the contexts, convincingly demonstrated the opposite effect. Preexposure to maze cues facilitated a discrimination involving the arms of the maze, but only when the preexposure and training were conducted in the same context.

To the extent that a perceptual learning effect depends on differential latent inhibition accrued to the elements, with more accruing to the common elements than the unique ones, one would expect results such as those reported by Trobalon et al. (1992). A change in context should cause some attenuation of latent inhibition, which would allow the common elements to become associated with the outcome, promoting generalization between the stimuli.

We did observe some latent inhibition in Group Intermixed during conditioning. To further assess the role of that latent inhibition in both the perceptual learning effect

we obtained in the previous experiments and its susceptibility to context, we conducted Experiment 2. Experiment 2 was designed to replicate the previous studies and included conditions that received a context switch between preexposure and conditioning, as did those in Trobalon et al. (1992).

Eight groups of participants were run in a 2 × 2 × 2 design manipulating the type of preexposure as well as the conditioning and test contexts. The design is shown at the bottom of Table 1. In the preexposure phase, participants played the game in Context A, in which they received intermixed preexposure to AX and BX, or no exposure. In the latter condition, they were simply exposed to the contextual stimuli as they were in the first experiments. Then, as indicated by the crossed lines in Table 1, participants from each group received conditioning in either the same (Context A) or different (Context B) context as that in which preexposure occurred. Participants from each of these conditions were then tested with BX in either the same or different context as that in which the preexposure phase was conducted.

Method

Participants and Apparatus

Seventy-one participants with the same characteristics as those participating in the previous studies participated in Experiment 2. The apparatus was the same as that used in the previous studies.

Procedure

All nonexperimental interactions with participants were the same as those in the previous studies. Participants belonged to one of eight conditions that were referenced by a 3-letter sequence so that the first letter (I or N) referred to whether participants received intermixed preexposure or none. Two letters following a hyphen (S or D) referred to whether the contexts of conditioning and testing were the same as or different from the context used in the exposure phase. For example, in the following test condition, I-SD received intermixed preexposure, conditioning in the same context as preexposure, and testing in a different context.

As in Experiments 1A and 1B, participants began by playing the game for 60 sec in a context different from that in which the preexposure phase would occur to ensure familiarity with that context. Context switches were initiated, as in Experiments 1A and 1B. Participants then received intermixed preexposure to the AX and BX stimuli, or none at all, in exactly the same manner as did the corresponding groups in the previous studies. For participants who received conditioning in the same context as preexposure (conditions I-SS, I-SD, N-SS, N-SD), conditioning of AX occurred as it did in the previous studies. For participants receiving conditioning in a context different from preexposure (I-DS, I-DD, N-DS, N-DD), a context switch was initiated. Following conditioning, testing of BX was conducted as in the previous experiment, with participants in the conditions I-SD, I-DS, N-SD, and N-DS receiving a context switch immediately prior to the test.

Results

Exclusion Criteria

There were 10 participants in each of the I- conditions—except for I-SD, which had 9. There were 8 in each of the N- conditions. Application of the exclusion criteria eliminated 2 participants from each of conditions I-SS, I-DS, N-DS, and I-DD, 3 participants from condition I-SD, and 1 participant from condition N-SD.

Exclusion was independent of condition ($\chi^2 = 5.67$, $p = .58$).

Conditioning

Suppression ratios. The data from conditioning are shown at the left in Figure 3. That portion of the figure collapses across the test context variable. The 10 conditioning trials were analyzed with an exposure × conditioning context (same as or different from that in preexposure) × test context (same as or different from preexposure) × trials ANOVA. There was an effect of trials [$F(9,459) = 53.79, p < .001$], and an exposure × conditioning context × trials interaction [$F(9,459) = 2.67, p = .004$]. There were no other effects in the overall analysis ($ps > .07$). Simple effect tests to investigate the interaction showed that acquisition of suppression was slower among participants who received intermixed preexposure than those who received none, yet the effect appeared at slightly different places, depending on where conditioning was conducted. When conditioning was conducted in the same context as preexposure, significant differences between those preexposed and those not were observed on Trials 2 and 3 [$Fs(1,358) > 5.74, ps < .02$]. When conditioning took place in the context different from preexposure, acquisition was still slower in the participants receiving intermixed exposure than those receiving none, but the differences occurred on Trials 4, 5, and 6 [$Fs(1,358) > 4.70, ps < .03$].

Pre-CS. The same analysis was applied to responding during the pre-CS. The analysis showed an effect of trials [$F(9,459) = 2.81, p = .003$] and some unexpected results whose pattern appeared spurious and was not simple to describe. There was a test context × trials interaction [$F(9,459) = 5.35, p < .001$] and an exposure × test context × trials interaction [$F(9,459) = 2.06, p = .03$]. Pre-CS response rates averaged 17.38 ($SD = 5.52$) at the beginning of training and 18.06 ($SD = 6.57$) at the end. Analyses to explore the interaction showed that in the intermixed conditions, there was slightly higher responding during the last four trials in the groups to be tested in the different context ($M = 21.37, SD = 4.12$) than those to be tested in the same context ($M = 14.89, SD = 6.92$) [$Fs(1,102) \geq 5.26, ps \leq .02$]. There were no such differences in the groups receiving no preexposure ($Fs < 1$). Responding averaged 18.96 ($SD = 7.27$) and 19.05 ($SD = 5.91$) in the participants to be tested in contexts same and different, respectively. There were no other effects in the overall analysis ($ps > .18$).

Because pre-CS response rate differences could complicate interpretation of the analysis of the suppression ratios, we also conducted the same analysis of the number of responses during the CS. In accordance with the analysis of the suppression ratios, there were no effects involving the test context variable ($ps > .10$). There was an effect of trials [$F(9,459) = 47.17, p < .001$], but the exposure × conditioning context × trials interac­tion seen in the suppression ratio data only approached significance [$F(9,459) = 1.81, p = .06$]. There were no other effects ($ps > .10$). Thus, whether the three-way
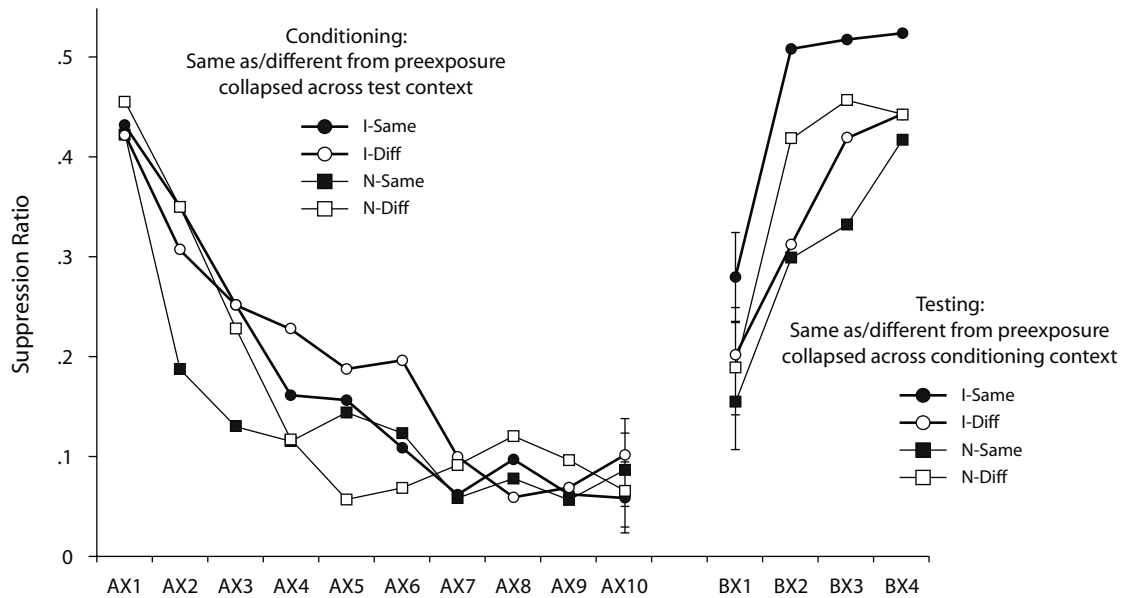
Figure 3. Data are shown from Experiment 2. The conditioning data shown at left collapse across the test context variable, and the test data shown at right collapse across the conditioning context variable. Left shows acquisition of suppression to AX following intermixed (I) preexposure to AX and BX, or none (N). Acquisition was conducted in either the same context as preexposure (solid symbols) or a different context (open symbols). Generalization testing with BX is shown at the right with testing conducted in either the same context as preexposure (closed symbols) or a different context (open symbols).

interaction observed with the suppression ratio data was the result of differences in responding controlled by the CS is suspect.

Test

   Suppression ratios. The test trials were analyzed with an exposure $\times$ conditioning context (same as or different from preexposure) $\times$ test context (same as or different from preexposure) $\times$ trials ANOVA. That analysis showed a reliable effect of trials [$F(3,153) = 42.34$, $p < .001$] and an exposure $\times$ test context interaction [$F(1,51) = 7.34, p = .009$]. There were no other effects in the overall analysis ($ps > .09$). Of importance, no effect of conditioning context was observed [$F(1,51) = 1.41$, $p = .21$], and it did not interact with any other factor in the analysis ($Fs < 1$).

   The lack of a three-way interaction suggested comparable exposure $\times$ test context interactions within each level of conditioning context. Subsequent exposure $\times$ test context analyses within each of the two levels of the conditioning context variable confirmed that interpretation. Exposure $\times$ test context interactions of identical form were observed regardless of where conditioning occurred [$Fs(1,51) \geq 9.17, ps < .004$]. Data are shown at the right in Figure 3, which collapses across the conditioning context variable.

   Simple-effect tests (all Fs, with 1 and 51 df) collapsed across the conditioning context factor confirmed that participants receiving intermixed preexposure showed less suppression to BX than did participants receiving none when testing was in the same context as was used in the preexposure phase ($F = 9.33, p = .004$). Participants re-

ceiving intermixed preexposure also showed less suppression than participants receiving intermixed preexposure with testing in a context different than preexposure ($F = 4.92, p = .03$). There were no other differences among the groups ($Fs < 2.58, p \geq .11$).

   Pre-CS. The same set of analyses above were applied to responding during the pre-CS and revealed no effects on the test ($ps > .10$). Responding during the pre-CS averaged 17.94 ($SD = 5.73$) during the test.

Discussion

   In the present experiment, participants received intermixed preexposure to the conditioning and test stimuli, or to none. Conditioning was subsequently conducted in either the same context as that used during the preexposure phase, or in a different context. The context of testing was factorially combined with the context of training so that testing was also conducted in the same or different context as that used in the preexposure phase. The results of the test replicated the previous study. Intermixed preexposure to the stimuli substantially reduced generalization, but only when testing occurred in the same context as preexposure. This reduction was true regardless of where conditioning took place. The critical variable was the match between the context of testing and the context in which perceptual learning took place.

   The results conceptually replicate those of Trobalon et al. (1992), who showed that a change of context between preexposure and conditioning attenuated perceptual learning. Here, participants that received preexposure in one context and conditioning and testing in another showed strong generalization. Participants that received preexpo-

sure, conditioning, and testing in the same context showed reduced generalization.

Conditioning in a context different from that used for preexposure reduced the effects of that preexposure on the test. Yet, that reduction was limited to participants who were also tested in the different context from that in which the preexposure phase was conducted. After conditioning occurred in a context different from preexposure, testing back in the preexposure context appeared to restore the effects of preexposure. That result was consistent with what is empirically known about latent inhibition. Although the change in context between preexposure and conditioning should have allowed excitation to accrue to the common elements, excitation acquired after latent inhibition has been shown to be context specific (Swartzentruber & Bouton, 1992; Westbrook, Jones, Bailey, & Harris, 2000). Thus, a change out of the conditioning context should attenuate responding controlled by the common elements and lead to less apparent generalization between the stimuli.

The results obtained with the final condition that received intermixed preexposure do not fit with the explanations offered above. As did participants in the first experiment, these participants received preexposure and conditioning in one context, and testing in another. To the extent that latent inhibition prevents the acquisition of associatively controlled responding to the common X elements during conditioning, a change in context after conditioning would be predicted to have little impact after conditioning. If the preexposure to the common elements leads to context specificity of subsequent learning, as was mentioned earlier, then a change in context should attenuate any responding controlled by those elements, leading to a further reduction in generalization rather than the increase in response suppression, which was observed.

The results from the conditioning phase are likewise only partially consistent with an explanation of perceptual learning on the basis of what is empirically known about latent inhibition. To the extent that latent inhibition accrued during preexposure, slower conditioning should have been observed in the groups that received intermixed preexposure as compared with those that received none. The pattern of results was generally in accord with that expectation. To the extent that latent inhibition is context specific with this method (Nelson & Sanjuan, 2006), we would expect to see the difference between the groups only when the conditioning and preexposure contexts were the same, and we did not observe such an effect. There were comparable effects of preexposure in each context during conditioning. Any latent inhibition accrued to these stimuli was not affected much by a context change, yet the mechanisms responsible for the reduction in generalization were. Thus, although latent inhibition appeared to be present, these results could be taken as evidence that latent inhibition had little to do with either the enhanced discrimination established by the intermixed preexposure procedure with this method, or its loss when the contexts of preexposure and testing differed.

The overall pattern of results across conditioning is consistent with the theory of McLaren and Mackintosh (2000). Like that of Wagner (1981), their theory assumes that when a stimulus is presented, it enters into associations with the context and other stimuli. Thus, it can be said that the stimulus is expected in that particular context and its associability is therefore reduced, resulting in a latent-inhibition effect. In McLaren and Mackintosh's terms, the salience of the stimulus is reduced by these interstimulus associations. When the stimulus is presented in a new context, its salience is restored, and conditioning can proceed more normally. McLaren and Mackintosh's variation of this idea provides an explanation for why context specificity of latent inhibition was not observed in the present study (see McLaren & Mackintosh, 2000, pp. 222–224). The stimuli presented were complex, being deliberately created from at least two discriminable components. Presented together, these components should become associated with each other. Establishment of these associations would compete with the establishment of associations between the context and the elements. Thus, when AX is presented in a new context, X would still be expected because of A, and vice versa, allowing the reduction in their salience to be maintained in the new environment.

McLaren and Mackintosh's (2000) ideas regarding how salience modulation occurs might also account for the test results observed with BX in the groups that received intermixed preexposure. When tested with BX, A was not present; thus, the effective salience of X should be restored, allowing it to produce more conditioned responding. The removal of the influence of A might also allow for the opportunity to better observe the effects of any associations between X and the context. Thus, when the context of testing was different from that in which preexposure took place, X could be even more effectively salient and bring about more of a response.

The explanation offered above fares somewhat less well in the conditions that did not receive preexposure. Here, any associations between X and the context could be formed only during the conditioning phase. Thus, we would have expected that any change out of the conditioning context in these groups should restore the salience of X and produce more suppression. No such effects were observed, and the trends were in a direction opposite of one favoring that hypothesis.

GENERAL DISCUSSION

The present studies were undertaken to observe whether a context switch could create a flattening of the generalization gradient. A perceptual learning manipulation was included in half of the conditions to enhance the discrimination between similar stimuli to allow for more room to observe an effect. Compound stimuli AX and BX were used in a video-game task. When participants learned to suppress their ongoing behavior in the presence of AX, substantial generalization was observed to BX. Simple intermixed preexposure to the stimuli reduced that generalization.

When the stimuli were tested in a context that was different from that in which conditioning took place, there was a small effect of context on conditioning, in that suppression to AX was less in the different context. Such an

effect could reflect a generalization decrement in the perception of the CS, or some type of hierarchical contextual control of the stimulus–outcome association by the context (see Holland, 1992). Regardless of the mechanism, the behavioral result reflects one form of flattening of the generalization gradient in which responding to the trained stimulus is decreased.

Despite the small loss of conditioning to the AX stimulus, there was a marked increase in generalization to BX, but only when that generalization had been reduced by intermixed preexposure. Such an increase represents the other form of flattening of the generalization gradient in which responding in the tails of the gradient increases.

The effect of a context change on the BX stimulus was restricted to groups that underwent perceptual learning. Perceptual learning is known to involve latent inhibition (see Hall, 1991; McLaren & Mackintosh, 2000). Thus, we extended the design of Experiment 1 with Experiment 2 to assess whether variations in the susceptibility of latent inhibition to contextual control might account for the results we obtained. As was discussed in the discussion of Experiment 2, there were aspects of the results that were consistent with both empirical facts and theories regarding latent inhibition, but neither provided a full account.

The effects of a context change on generalization observed here are anticipated by at least two other theories of perceptual learning (Petrov et al., 2006; Saksida, 1999). These models instantiate perceptual learning in neural networks as the result of competitive weight adjustment (Saksida, 1999) or Hebbian reweighting (Petrov et al., 2006). The model of Saksida (1999) assumes that the elements of a stimulus converge their input on intermediate units that are then associated with responses. The convergence is driven by a competitive learning algorithm based on the distance between the intermediate unit and the input. The result of inputting similar stimuli to such a network over trials would be the modification of the connection weights so that the units activated in the intermediate layer are further apart than they are in the input layer.

The model of Petrov et al. (2006) is considerably more complex, although that detail need not be considered here. The relevant aspects are that a pattern of input is passed into a representational subsystem that is then connected to a decision unit (e.g., pattern oriented right or left, same or different) via an adjustable weight matrix.

A change of context will affect both of these models for essentially the same straightforward reason. The novel input combination produced by the stimulus in the new context will produce an activation of a set of weights in the respective matrices that is suboptimum in comparison with their state at the end of training. Thus, units within the systems that have come to be activated by a particular pattern of input would be less active, and units that have come to be inactive given a particular pattern of input would become somewhat more active. In this way, similarity between inputs would be achieved. In the model of Saksida (1999), there would be less discrete representations of different inputs at the competitive layer, and in that of Petrov et al. (2006), there would be a less precise activation of the decision unit. In both cases, the behavioral result should mimic the present results. An increased similarity between AX and BX—whether it is represented in a competitive layer of intermediary units or in the output of a decision unit—would result in an increase in suppression to BX and a decrease in suppression to AX.

As was elaborated in the introduction, flattening of a generalization gradient appears to occur when the stimuli are likely to be perceived as similar. When the gradient may be sharp due to physical differences between stimuli, flattening is not observed (see, e.g., Zhou & Riccio, 1996). Yet, in the present studies, when the gradient was sharp due to perceptual learning, a change in context had a substantial effect. These results raise the possibility that when flattening of generalization gradients is observed, some perceptual learning may have been involved. When exposure to stimuli leads to perceptual learning, enhancing an otherwise broad gradient, flattening may be especially marked. Experiments in which changes in generalization over time are clear and less open to alternative explanations used methods that involved multiple trials with stimuli that could be considered relatively complex, such as goal boxes (Perkins & Weyant, 1958), shock and escape compartments in avoidance learning (McAllister & McAllister, 1963), and line drawings of teacups (Bahrick, Clark, & Bahrick, 1967). Multiple exposures to complex stimuli over trials, as in the work of Petrov et al. (2006), are the conditions that should allow for the operation of a perceptual learning mechanism.

In studies that assessed the effects of time on flattening generalization gradients, there are examples in which either the training stimuli (see, e.g., Wiltgen & Silva, 2007) or both the training and test stimuli were preexposed (Desiderato et al., 1966; Rosas & Bouton, 1997). These studies showed marked effects of time on generalization. For example, Rosas and Bouton (1997) took advantage of the context specificity of latent inhibition to assess rats' ability to discriminate contexts. The animals received equal exposure to the training and test contexts to ensure equal familiarity with those stimuli. When conditioning of a tone CS was conducted immediately after its preexposure, the animals could readily discriminate the context. Latent inhibition to the tone was observed when it was conditioned in the context in which latent inhibition was acquired during preexposure to the tone, but not when the tone was conditioned in the different context. However, animals appeared to be unable to discriminate the contexts 28 days later. Latent inhibition was equally present in the same and different contexts. Although suggestive, these studies lacked controls that did not receive such preexposure in order to allow a full assessment of whether the preexposure to the contextual stimuli contributed to their effect.

The present research has shown that in the absence of any explicit perceptual learning with both training and test stimuli, a change in the context results in a small loss of responding to the training stimulus, with no observable effect on the test stimulus. Such a result characterizes one form of flattening of the generalization gradient. After perceptual learning, when the generalization gradient has been sharpened, the flattening can be observed more apparently with a change in context. A change in context

leads to a small loss of responding to the training stimulus and to a marked increase in responding to the test stimulus. The magnitude of the effect of time—sometimes considered a form of context (e.g., Bouton et al., 1999)—on the flattening of a generalization gradient could partly depend on the perceptual learning that occurs during the participant's experimental and preexperimental history with the elements of the stimuli in question.

AUTHOR NOTE

REFERENCES

BAHRICK, H. P., CLARK, S., & BAHRICK, P. (1967). Generalization gradients as indicants of learning and retention of a recognition task. *Journal of Experimental Psychology*, **75**, 464-471.

BOUTON, M. E., NELSON, J. B., & ROSAS, J. M. (1999). Stimulus generalization, context change, and forgetting. *Psychological Bulletin*, **125**, 171-186.

BUTLER, B. M. (1964). Postconditioning delay and stimulus similarity as factors in the measurement of acquired fear. *Connecticut College Psychology Journal*, **1**, 3-18.

CHANNELL, S., & HALL, G. (1981). Facilitation and retardation of discrimination learning after exposure to the stimuli. *Journal of Experimental Psychology: Animal Behavior Processes*, **7**, 437-446.

CRIST, R. E., KAPADIA, M. K., WESTHEIMER, G., & GILBERT, C. D. (1997). Perceptual learning of spatial localization: Specificity for orientation, position, and context. *Journal of Neurophysiology*, **78**, 2889-2894.

DESIDERATO, O., BUTLER, B. [M.], & MEYER, C. (1966). Changes in fear generalization gradients as a function of delayed testing. *Journal of Experimental Psychology*, **72**, 678-682.

GOLDSTONE, R. L. (1998). Perceptual learning. *Annual Review of Psychology*, **49**, 585-612.

HALL, G. (1991). *Perceptual and associative learning*. Oxford: Oxford University Press, Clarendon Press.

HALL, G. (2003). Learned changes in the sensitivity of stimulus representations: Associative and nonassociative mechanisms. *Quarterly Journal of Experimental Psychology*, **56B**, 43-55.

HOLLAND, P. C. (1992). Occasion setting in Pavlovian conditioning. In D. L. Medin (Ed.), *The psychology of learning and motivation* (Vol. 28, pp. 69-125). New York: Academic Press.

KRAEMER, P. J. (1984). Forgetting of visual discriminations by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, **10**, 530-542.

LAVIS, Y., & MITCHELL, C. (2006). Effects of preexposure on stimulus discrimination: An investigation of the mechanisms responsible for human perceptual learning. *Quarterly Journal of Experimental Psychology*, **59**, 2083-2101.

LUBOW, R. E. (1989). *Latent inhibition and conditioned attention theory*. New York: Cambridge University Press.

LUBOW, R. E., RIFKIN, B., & ALEK, M. (1976). The context effect: The relationship between stimulus preexposure and environmental pre-exposure determines subsequent learning. *Journal of Experimental Psychology: Animal Behavior Processes*, **2**, 38-47.

McALLISTER, W. R., & McALLISTER, D. E. (1963). Increase over time in the stimulus generalization of acquired fear. *Journal of Experimental Psychology*, **65**, 576-582.

McLAREN, I. P. L., & MACKINTOSH, N. J. (2000). An elemental model of associative learning: I. Latent inhibition and perceptual learning. *Animal Learning & Behavior*, **28**, 211-246.

NELSON, J. B., & SANJUAN, M. D. C. (2006). A context-specific latent inhibition effect in a human conditioned suppression task. *Quarterly Journal of Experimental Psychology*, **59**, 1003-1020.

PERKINS, C. C., JR., & WEYANT, R. G. (1958). The interval between training and test trials as a determiner of the slope of generalization gradients. *Journal of Comparative & Physiological Psychology*, **51**, 596-600.

PETROV, A. A., DOSHER, B. A., & LU, Z.-L. (2006). Perceptual learning without feedback in non-stationary contexts: Data and model. *Vision Research*, **46**, 3177-3197.

RICCIO, D. C., RICHARDSON, R., & EBNER, D. L. (1984). Memory retrieval deficits based upon altered contextual cues: A paradox. *Psychological Bulletin*, **96**, 152-165.

ROSAS, J. M., & BOUTON, M. E. (1997). Additivity of the effects of retention interval and context change on latent inhibition: Toward resolution of the context forgetting paradox. *Journal of Experimental Psychology: Animal Behavior Processes*, **23**, 283-294.

SAKSIDA, L. M. (1999). Effects of similarity and experience on discrimination learning: A nonassociative connectionist model of perceptual learning. *Journal of Experimental Psychology: Animal Behavior Processes*, **25**, 308-323.

SATTERTHWAITE, F. E. (1946). An approximate distribution of the estimates of variance components. *Biometrics Bulletin*, **2**, 110-114.

SWARTZENTRUBER, D., & BOUTON, M. E. (1992). Context sensitivity of conditioned suppression following preexposure to the conditioned stimulus. *Animal Learning & Behavior*, **20**, 97-103.

SYMONDS, M., & HALL, G. (1995). Perceptual learning in flavor aversion conditioning: Roles of stimulus comparison and latent inhibition of common stimulus elements. *Learning & Motivation*, **26**, 203-219.

THOMAS, D. R., WINDELL, B. T., BAKKE, I., KREYE, J., KIMOSE, E., & APOSHYAN, H. (1985). Long-term memory in pigeons: I. The role of discrimination problem difficulty assessed by reacquisition measures. II. The role of stimulus modality assessed by generalization slope. *Learning & Motivation*, **16**, 464-477.

TROBALON, J. B., CHAMIZO, V. D., & MACKINTOSH, N. J. (1992). Role of context in perceptual learning in maze discriminations. *Quarterly Journal of Experimental Psychology*, **44B**, 57-73.

WAGNER, A. R. (1981). SOP: A model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 5-47). Hillsdale, NJ: Erlbaum.

WELCH, B. L. (1938). The significance of the difference between two means when the population variances are unequal. *Biometrika*, **29**, 350-362.

WESTBROOK, R. F., JONES, M. L., BAILEY, G. K., & HARRIS, J. A. (2000). Contextual control over conditioned responding in a latent inhibition paradigm. *Journal of Experimental Psychology: Animal Behavior Processes*, **26**, 157-173.

WILTGEN, B. J., & SILVA, A. J. (2007). Memory for context becomes less specific with time. *Learning & Memory*, **14**, 313-317.

ZHOU, Y., & RICCIO, D. C. (1996). Manipulation of components of context: The context shift effect and forgetting of stimulus attributes. *Learning & Motivation*, **27**, 400-407.